

Task-Based Algorithm for Matrix Multiplication: A Step Towards Block-Sparse Tensor Computing

Justus A. Calvin
justusc@vt.edu

Edward F. Valeev
efv@vt.edu

Department of Chemistry
Virginia Tech
Blacksburg, VA

ABSTRACT

Distributed-memory matrix multiplication (MM) is a key element of algorithms in many domains (machine learning, quantum physics). Conventional algorithms for dense MM rely on regular/uniform data decomposition to ensure load balance. These traits conflict with the irregular structure (block-sparse or rank-sparse within blocks) that is increasingly relevant for fast methods in quantum physics. To deal with such irregular data we present a new MM algorithm based on Scalable Universal Matrix Multiplication Algorithm (SUMMA). The novel features are: (1) multiple-issue scheduling of SUMMA iterations, and (2) fine-grained task-based formulation. The latter eliminates the need for explicit internodal synchronization; with multiple-iteration scheduling this allows load imbalance due to nonuniform matrix structure. For square MM with uniform and nonuniform block sizes (the latter simulates matrices with general irregular structure) we found excellent performance in weak and strong-scaling regimes, on commodity and high-end hardware.

Keywords

matrix multiplication, tensor contraction, task parallelism, SUMMA, latency tolerance, distributed computation

1. INTRODUCTION

High performance parallel algorithms for matrix multiplication (MM) — the most important special case of the general tensor contraction, and often its building block — have been studied for decades [13, 22, 8, 32, 28, 17, 50, 23]. MM nevertheless continues to draw attention of researchers [46, 49, 5, 19, 31, 37, 47, 40, 41, 42] due to the continuing evolution of the computer hardware as well as the prominent role of matrix and tensor computation in a variety of scientific domains, such as physics of classical and quantum fields (most notably, electronic structure [7, 25, 20, 47]) as well as data analysis and model building in ma-

chine learning [29], chemometrics [44], neuroscience [33], and many more. A frontier challenge posed by the scientific domain needs is the increasing importance of sparse and non-standard tensorial data representations (dense and sparse tensors, multiresolution spectral element trees, H matrices, matrix product states, tensor networks, *etc.*). Another major challenge is the increasing complexity of communication hierarchy and the continuing increase of the communication cost, relative to that of computation. This spurs the search for algorithms that minimize/avoid communication and/or hide its cost. Here we explore improvements of a standard dense matrix-multiplication algorithm that can hide communication costs and tolerate network latencies, data inhomogeneity, and sparsity, which serves as a platform for the development of algorithms on irregular tensorial data structures.

To demonstrate the tension between standard dense MM algorithms and the needs of emerging scientific domains, consider a concrete example of chemistry and materials science. The matrices in such context represent quantum states (of electrons) and operators represented in some basis. Efficient application of operators to states — represented by matrix multiplication — demands taking advantage of the matrix structure arising from the physical properties of the entities these matrices represent. Such structure could be (a) block-sparsity due to the distance decay of the operator kernel and the localized nature of basis functions [26], (b) symmetries under geometric and other transformations [41, 25], or (c) block-rank-sparsity due to near-sightedness of physical interactions [35, 34, 38] (*i.e.* blocks of the matrix/tensor are dense but have low-rank representations; in the applied math community related matrix structures are known as H matrices). Notably, the matrix structure is affiliated with the problem-specific blocking of matrix dimensions that arises due to domain-specific needs and often cannot be chosen arbitrarily. In other words, the matrices that we encounter are “sparse” in a general sense, which encompasses element, block, and block-level-rank sparsity; but in a practical sense the matrices are not sparse enough to be a good match for sparse MM algorithms. Nonuniform blocking and inhomogeneity of data due to the matrix structure conflict with the *uniform* data distribution of standard distributed-memory dense MM algorithms — including Cannon’s [13], SUMMA [50], and others [22, 17, 10, 47, 41].

To design algorithms for multiplication of matrices with an irregular structure, we propose a reformulation of standard, dense MM algorithms that allows them to be tolerant of data inhomogeneity by expressing them in terms of

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SC15, 27th International Conference for High Performance Computing, Networking, Storage and Analysis Austin, TX, USA

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$15.00.

tasks. By using the standard algorithm as the framework our approach will be potentially optimal for dense uniformly-blocked matrices, and yet capable of handling irregularity in the matrix structure. Task-based/dataflow programming models are a natural choice for implementation of algorithms with irregular data and computation patterns; such models have been used successfully for dense matrix algebra applications [27, 9]. Besides handling matrices with structure, a task based approach will provide additional benefits: (a) inter-node communication costs can be partially or fully hidden by overlapping computation and communication, (b) performance should be less sensitive to topology and tolerant of latency and CPU clock variations, (c) fine-grained task-based parallelism is a proven means to attain high intra-node performance by leveraging massively multicore platforms and hiding the costs of memory hierarchy (*e.g.* Intel TBB, Cilk), (d) lack of global synchronization allows the overlap multiple high-level stages of computation (*e.g.* multiple matrix multiplications contributing to the same expression).

This paper describes our approach, its implementation, and performance for matrix multiplication of square matrices with uniform and nonuniform blocking. Our implementation is available as part of the open-source project TiledArray (a general-purpose tensor library) [12].

2. DISTRIBUTED-MEMORY MATRIX MULTIPLICATION

This section contains an overview of the relevant literature on dense matrix multiplication algorithms for distributed memory computers. Our discussion does not consider *fast* algorithms [18], *e.g.* Strassen [48], that have computational complexity lower than $\Theta(N^3)$ of the naïve algorithm, where N is the matrix size. Thus for our purposes, all algorithms are computation-optimal and only differ in the algorithm construction, communication patterns, and memory usage per node.

An early parallel matrix-multiplication algorithm that has gained widespread use is Cannon’s algorithm. In this method, the input and output matrices are embedded onto a two-dimensional (2D) mesh of P processes (nodes), where the input matrices are moved via a series of row and column rotations in systolic loops [13]. Cannon’s algorithm asymptotically meets the lower bound on communication, $\Omega(N^2/\sqrt{P})$ [30], for 2D algorithms that require $\mathcal{O}(N^2/P)$ memory per process. The broadcast-multiply-roll algorithm of Fox *et al.*, which is also asymptotically optimal, introduced the use of a (pipelined) broadcast for movement of one of the two input matrices [22, 21]. Unfortunately neither algorithm is universally applicable [50, 43]. Some of the limitations were removed in the Transpose algorithm of Lin and Snyder [32], and further by Choi and coworkers in Parallel Universal Matrix Multiplication Algorithm (PUMMA) [17] and by Huss-Lederman *et al.* in the BiMMer algorithm [28]. Perhaps the simplest and easiest to generalize is a group of 2D algorithms utilizing row/column broadcasts followed by rank- k updates. These algorithms were developed independently by several groups [2, 16, 50], most notably by van de Geijn and Watts who dubbed the approach Scalable Universal Matrix Multiplication Algorithm (SUMMA) [50]. The simplicity and flexibility of SUMMA made it very successful (*e.g.* it is the standard implementation of GEMM in ScaLAPACK), and

has motivated several variants [14, 23, 10]. In addition, the core ideas are also widely used in linear algebra algorithms [39, 46, 47, 41]. Since our work is based on SUMMA, it is described in more detail in Section 2.1.

All 2D algorithms described so far are optimal with respect to memory use. However, it is possible to reduce the communication cost further by using a three-dimensional (3D) mesh of processes at the cost of additional memory usage [30, 4]. In 3D algorithms [8, 1], each matrix is replicated across one dimension of a $\sqrt[3]{P} \times \sqrt[3]{P} \times \sqrt[3]{P}$ mesh. This data distribution reduces the communication cost by a factor of $P^{1/6}$ [4], but requires $\mathcal{O}(N^2/P^{2/3})$ data per node, a factor of $P^{1/3}$ increase over the 2D case. The gap between the 2D and 3D algorithms is bridged by the so-called 2.5D algorithms [46] that are communication-optimal with and without memory constraints [4].

Although 2D MM algorithms are not universally communication-optimal, we view them as key for two reasons: (a) they are the only choice under tight memory constraints, which is often the case in data intensive applications, and (b) they are a building block for higher-dimensional MM algorithms [46]. Thus in our work we focus on SUMMA as the most popular 2D algorithm.

2.1 2D SUMMA and Its Variants

SUMMA implements matrix multiplication $\mathbf{C} = \mathbf{AB}$ as a series of rank- k updates. The input and output matrices are embedded on a rectangular process mesh in an element-cyclic or block-cyclic manner to ensure approximate load balance. In each iteration of the algorithm a column/row panel of \mathbf{A}/\mathbf{B} is broadcast along rows/columns of the process grid, respectively; matrix \mathbf{C} remains stationary throughout the procedure (variants of SUMMA in which \mathbf{A} or \mathbf{B} are stationary are also possible; transposed multiplies, *e.g.* $\mathbf{C} = \mathbf{AB}^\top$, are also relatively simple to handle [50, 43]). Each pair of broadcasts is followed by a rank- k update, $C_{ij} \leftarrow A_{ik}B_{kj} + C_{ij}$, (Einstein summation convention is used throughout). Original SUMMA papers by van de Geijn and Watts [50] and by Agarwal *et al.* [2] considered versions of the algorithm that overlapped computation and communication by pipelining and preemptive broadcasts, respectively, and other broadcast variants have been considered [43], including topology-specific broadcasts [45]. For simplicity, we present a SUMMA version with preemptive broadcasts in Figure 1.

Algorithm 1 SUMMA with non-blocking broadcast

```

1: BROADCAST( $A(*, 0), 0, row\_group$ )
2: BROADCAST( $B(0, *), 0, col\_group$ )
3: for  $k = 0, \dots, K - 1$  do
4:   if  $k + 1 < K$  then
5:      $row\_root \leftarrow (k + 1) \bmod cols$ 
6:      $col\_root \leftarrow (k + 1) \bmod rows$ 
7:     BROADCAST( $A(*, k + 1), row\_root, row\_group$ )
8:     BROADCAST( $B(k + 1, *), col\_root, col\_group$ )
9:   end if
10:  WAIT( $A(*, k)$ )
11:  WAIT( $B(k, *)$ )
12:   $C(*, *) \leftarrow \alpha A(*, k) \cdot B(k, *) + C(*, *)$ 
13: end for

```

DIMMA, an early variation of SUMMA introduced by

Choi in 1998, improved performance of synchronous SUMMA by reducing the slack in communication [14, 15]. The key insight of this work was that the order of broadcasts in the original SUMMA implementation, coupled with the communication barriers, created a significant slack in communication. Iterations were reordered in DIMMA such that each node broadcasts all of its data in succession as opposed to the round-robin approach in SUMMA. Similar improvements can be attained by overlapping communication and computation [2].

SUMMA was recently extended to sparse MM (SpSUMMA) by Buluç and Gilbert [10, 11]. This algorithm is similar to dense SUMMA, except matrix sub-blocks are stored in a doubly compressed sparse column (DCSC) format and sparse generalized matrix-multiplication (SpGEMM) is used to compute rank- k updates. The main challenges of all sparse MM algorithms, including SpSUMMA, are the increased relative costs of communication compared to the dense case, load imbalance, and the relatively low intranode performance of sparse matrix kernels.

The problem of load imbalance does not appear in dense SUMMA implementations as the work load is nearly-optimally balanced. With random sparsity approximate load balance is achieved in the asymptotic limit, however with structured sparsity (e.g. matrices with decay) one does not expect rows/columns to be uniformly filled.

In the regime of high sparsity (low matrix fill-in factors) the communication time dominates the computation time due to the effectively-reduced benefit of blocking. The increasing role of communication in sparse MM can be somewhat alleviated by communication hiding. Buluç and Gilbert discussed potential benefits of communication hiding for sparse MM in Ref. [10] but did not pursue this approach in their experiments due to lack of quality one-sided communication tools [11]. Another possibility is to minimize communication, e.g. by switching to 3D [3]. Although sparse MM 2D SUMMA is not strongly scalable, nevertheless good scalability of SpSUMMA in practice was demonstrated [11].

3. TASK-BASED 2D SUMMA

Our work to improve SUMMA is motivated by the needs of computation on matrices/tensors with irregular block structure in the context of many-body quantum physics. Some of the challenges of computing with such data are similar to the general challenges of sparse MM: increased communication/computation ratio and lack of load balance. The latter is compounded by the desire to use physics-based blocking of matrix dimensions. To address these challenges we decided to investigate a task-based formulation of SUMMA algorithm, to partially offset some communication costs and to alleviate the load imbalance. Prior efforts to reformulate dense matrix multiplication using a task-based programming models are known [9, 6]; the novelty of our effort is the focus on dense matrices/tensors as well as matrices/tensor with irregular structure, such as block-rank-sparsity, that cannot be handled straightforwardly using the standard dense-only approaches. In this section we describe the design of our algorithm, by highlighting the differences with the procedural SUMMA implementations[50] We first analyze the data dependencies of discrete operations in the procedural SUMMA implementation. We then discuss the task composition and dependencies of our modified implementation. For simplicity, we only consider the 2D SUMMA

implementation; our approach is immediately applicable to the 2.5D variant since it's based on 2D SUMMA.

3.1 Dependency Analysis of Standard SUMMA

Like all dense MM algorithms, SUMMA consists of data movement and computation tightly synchronized with each other (see Algorithm 1). Namely, the rank- k update, $C_{ij} \leftarrow A_{ik}B_{kj} + C_{ij}$, of each SUMMA iteration depends on the data from the broadcasts of the corresponding panels of **A** and **B** as well as the previous iteration's rank- k update. In addition to these *data dependencies*, each broadcast is synchronized with a prior rank- k update since communication operations are initiated at the beginning of each SUMMA iteration, as shown in Figure 1 (we denote such *sequence dependencies* by dashed lines). Such design ensures that only a minimal memory is used (technically, nonblocking broadcasts require more memory than optimal). However, such design limits the amount of parallelism available in a given iteration (see the Figure 1). Specifically, we can parallelize the rank- k update of C as well as the column and row broadcasts of A and B , but SUMMA iterations — although almost independent from one another — are executed serially, one after the other. Furthermore, such design is not tolerant of any source of load imbalance, due to, for example, processor clock variation, network congestion, or — most importantly for us — due to the inhomogeneity of data.

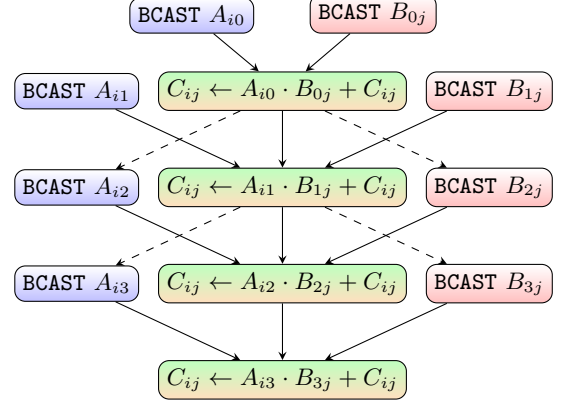


Figure 1: A directed acyclic graph of the procedural SUMMA implementation consisting of broadcast (BCAST) and rank- k updates. Solid edges indicate data dependencies, and dashed edges indicate sequence dependencies.

3.2 Task-Based Implementation

To tolerate the data inhomogeneity, whether due to block sparsity or block-rank sparsity, the task queue of each processes should have as many tasks as possible at any given time. To maximize the amount of available parallelism, we introduced the following modifications to the standard 2D SUMMA:

- data overdecomposition,
- multiple-issue schedule of SUMMA iterations, and
- task-decomposed broadcasts

Although the standard dense SUMMA allows near-perfect load balancing by using element-cyclic embedding of matri-

ces onto 2D mesh, we over decompose matrices into arbitrarily sized blocks, also embedded cyclically. In fact the dimension blocking is a central concept in TILEDARRAY library since it arises naturally in physical problems (hence the need to support arbitrary block sizes, including nonuniform blocking). These blocks should be sufficiently small to ensure that each process is assigned *many* blocks, yet large enough to perform acceptably with high performance BLAS3 library (used to perform block-level multiplies). Note that the overdecomposition also improves intra-node scalability when coupled with an optimized, parallel runtime like Intel TBB, and helps improve load balance in dealing with matrices/tensors that have permutational symmetries, e.g. in the CAST algorithm [41].

Decoupling of SUMMA iterations helps to increase the amount of parallelism beyond what is afforded by a single SUMMA iteration. This is similar to extracting more parallelism out of SUMMA by performing batches of SUMMA iterations in parallel on 2D subsets of a 3D process mesh in 2.5D MM algorithms [46], except here different iterations of 2D SUMMA will be concurrently executed and scheduled on the same 2D mesh. Note that the dependencies between SUMMA iterations occur via the result matrix \mathbf{C} . To remove these data dependencies we assume that there is enough memory available to split the rank- k update operation into two separate tasks: a matrix-multiplication task producing a temporary block, $C_{ij}^{(k)} = A_{ik} \cdot B_{kj}$, and a reduction of the temporary into the result, $C_{ij} = C_{ij}^{(k)} + C_{ij}$.

Figure 2 shows the directed acyclic graph (DAG) of task-based SUMMA with inter-iteration dependencies removed. By scheduling multiple iterations at once, we are able to overlap and interleave communication and computation of different iterations. Also note that, in principle, all iterations of SUMMA can potentially be scheduled concurrently. However, in practice we limit the number of *multiple-issue scheduling* to several iterations of SUMMA, and schedule additional iterations as the preceding iterations are retired, in a pipelined fashion. The additional task dependencies needed to implement this *throttling* mechanism are not shown in Figure 2. The number of concurrent iterations is determined by the dimensions of the process grid. For example, given a process grid with P_{row} rows and P_{col} columns, the number of concurrent iterations, I , is equal to,

$$I(P_{\text{row}}, P_{\text{col}}; K) = \begin{cases} 2, & \text{if } P_{\text{row}} < 2 \text{ or } P_{\text{col}} < 2 \\ K, & \text{if } P_{\text{row}} \geq K \text{ and } P_{\text{col}} \geq K \\ \min(P_{\text{row}}, P_{\text{col}}) \end{cases} \quad (1)$$

where K is the number of blocks along the inner dimension of the matrix multiplication. This limit is based on performance profiling, where we observed the maximum performance at $\min(P_{\text{row}}, P_{\text{col}})$.

Lastly, we decomposed broadcast tasks into subtasks each of which deals with a point-to-point communication. This achieves benefits similar to what pipelining does to standard 2D SUMMA [50, 49], by introducing the overlap between communication and computation of tasks working on a given SUMMA iteration. For example, a given process can start doing compute work as soon as the minimal amount of data — one block of \mathbf{A} and one block of \mathbf{B} — is available, concurrently with any additional communication of these blocks to other processes in the row and column groups. Work scheduling based on the availability of data in our

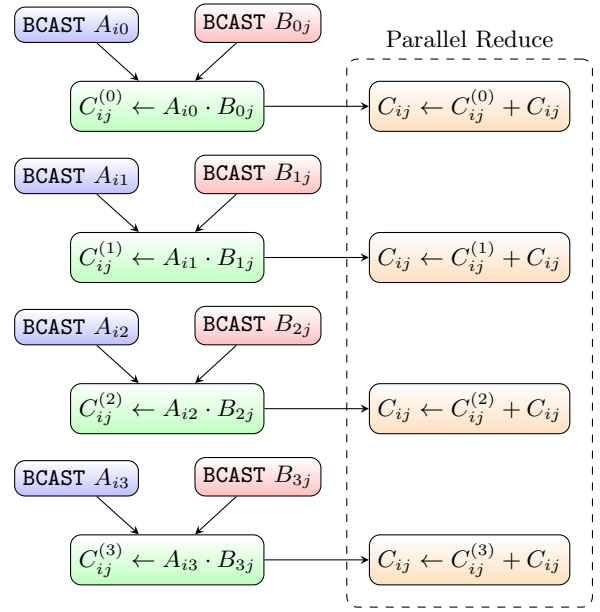


Figure 2: A direct translation of SUMMA to task-based implementation.

task approach should improve tolerance to latency and network topology [36]. In the current implementation we use a hardware-oblivious binary tree broadcast, *i.e.* it is not optimized for any particular interconnect hardware or topology, such as a multidimensional torus or fat tree. While our design does not preclude the use of optimized broadcast algorithms, a hardware oblivious algorithm should help performance portability across high-end and commodity hardware.

The dynamic scheduling of computation and communication trades off the predictable bounds on the resource use, in particular memory and network bandwidth, for high performance. For example, by multiple-issue schedule of SUMMA iterations increases memory use because each node must store an additional column/row block of the the argument matrices. Additional temporary storage might be used by matrix-multiplication producing temporaries subsequently reduced into the blocks of the target. The average memory overhead per node per SUMMA iteration, as depicted in Figure 2, is proportional to:

$$\underbrace{\frac{Mmk}{P_{\text{row}}}}_{\text{left-hand}} + \underbrace{\frac{Nkn}{P_{\text{col}}}}_{\text{right-hand}} + \underbrace{\frac{MNmn}{P_{\text{row}}P_{\text{col}}}}_{\text{result}} \quad (2)$$

where k is the average block size of the inner dimension; m and n are the average number of rows and columns in the result matrix blocks, respectively; M and N are the number of block rows and block columns in the result matrix, respectively; and P_{row} and P_{col} are the number of rows and columns in the process grid, respectively. However, we mitigate the memory overhead associated with the local result matrix by decomposing the rank- k -update tasks such that each matrix-multiplication task only computes a small sub-block of the local result matrix (see Figure 3). This decomposition of the rank- k -update and broadcast operations are similar to the decomposition of work between nodes within a SUMMA iteration, but without the analogous communica-

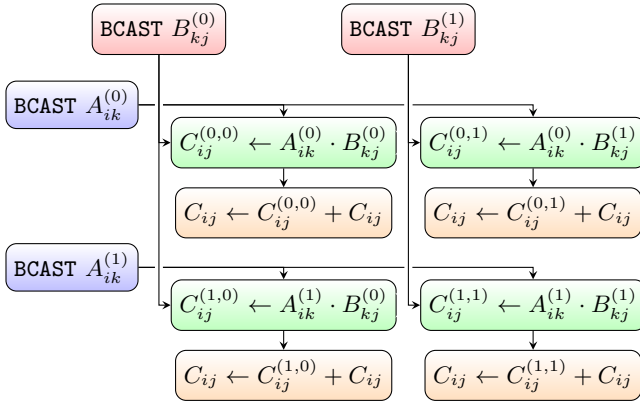


Figure 3: A directed acyclic graph for a single, decomposed SUMMA iteration. Broadcast, matrix multiply, and reduction tasks are decomposed such that each sub-block of the local result matrix, $C_{ij}^{(x,y)}$, depends on two sub-block broadcast tasks for $A_{ik}^{(x)}$ and $B_{kj}^{(y)}$.

tion (dependencies) between computation tasks. This allows different threads to work on different sub-blocks of the local result matrix concurrently. The maximum memory overhead is therefore bounded by the number of threads and only weakly dependent on the number of concurrent iterations. We further mitigate memory costs by combining the matrix-multiplication and reduction tasks whenever possible, which eliminates the memory and computational overhead of a separate reduction task. As a result, matrix-multiplication tasks only allocate additional, temporary memory for the result sub-blocks when two or more threads compute contributions to the same sub-block simultaneously. The memory cost associated with the left- and right-hand matrices is also mitigated by the overdecomposition of data. This is due to the fact that each row and column sub-block only remains in memory as long as it is needed, which is not necessarily for the entire duration of computation for an iteration. Finally, the overall memory consumption can be controlled by limiting the number of concurrent SUMMA iterations that are executed on each node. If I SUMMA iterations are scheduled at once each process’s memory potentially increases over that of standard SUMMA by a factor of I ; although as explained the resulting increase would in practice be much less than that.

4. RESULTS

The task-based 2D SUMMA algorithm is implemented in C++ in the TILEDARRAY library [12]. TILEDARRAY uses the MADNESS runtime [24] to express the computation in terms of tasks and to manage the low-level details of task scheduling and data movement. Both TILEDARRAY and MADNESS can be obtained under the terms of the GNU General Public License.

4.1 Experimental Setup

In our experiments, we compare the performance of the task-based SUMMA implementation for matrix multiplication of square uniformly- and nonuniformly-blocked real double-precision matrices; the square shape was chosen for simplic-

ity of the performance analysis and is not a constraint of our implementation. The Basic Linear Algebra Subprograms (BLAS) DGEMM routine was used as the block multiply-add kernel in our MM tasks. On the BG/Q system, we use the BLAS functions in IBM’s Engineering and Scientific Subroutine Library (ESSL). Intel’s Math Kernel Library (MKL) is used on the small, x86_64 Linux cluster. Although concurrent versions of DGEMM are available in both libraries, we use serial DGEMM to properly allocate resources with the MADNESS task-based parallel runtime system; this limitation will be removed in the near future. Reported wall times are an average of n repeated multiplications of two random matrices, where $n = 30$ for most computations except when run time limits constrained the sample sizes.

Block size for the uniform MM tests are optimized for each system, where we select the block size with the best performance from a series of single-process MM tests with a range of block sizes. To determine the size of blocks in the nonuniform MM tests, we first start by constructing M empty, row blocks, where M is equal to the number of row blocks in the uniformly-blocked matrices. We then randomly one to the size a row block, and repeat this step until the total number of rows among all blocks is equal to the number of rows in the uniformly blocked matrices. We repeat procedure for column block sizes. This ensures the average block size, as well as the number of blocks per matrix, in the nonuniform MM tests are equal to that of the uniformly-blocked matrices.

4.2 Performance of Task-Based SUMMA




Figure 4: Floating-point operation rate for multiplication of uniformly- and nonuniformly-blocked square matrices on IBM BG/Q. The matrix sizes included in the test are 32, 768, 65, 536, 98, 304, 256, 000.

Performance measurements on high-end hardware were performed on “Mira,” an IBM Blue Gene/Q (BG/Q) system at Argonne National Laboratory. For each matrix size we varied the number of compute nodes from 2^5 and 2^{14} (each node has 16 compute cores with 4 hardware threads




Figure 5: Wall time for multiplication of uniformly- and nonuniformly-blocked square matrices on IBM BG/Q. The matrix sizes included in the test are 32, 768, 65, 536, 98, 304, 256, 000.

per core; thus the largest computation utilized $2^{20} \approx 10^6$ threads).

In Figure 4 and 5, we see an approximately linear increase in computational rate for each matrix size. The combine data for all matrix sizes also shows approximately linear weak scaling in the entire range of processor counts. The performance difference between the uniform and nonuniform tests is small and follow each other closely, though the tests with uniform block sizes are consistently faster than those with nonuniform block sizes. In fact, slopes for both the uniform and nonuniform rate tests are approximately equal. This suggests that the performance difference between these two tests is due to an intra-node effect, rather than a scaling issue that one would expect to see with a significant load imbalance. This difference is an expected result since DGEMM performance is sensitive to the matrix size for small matrices.

We also perform a strong scaling test on a small, fifteen-node cluster with two eight-core 2.60 GHz Intel Xeon E5-2670 processors (16 cores total) and 16 GB of RAM per node, with an Infiniband switch. We use a matrix size of 32,768 for all tests on this system. The block size is set to 256. Similarly, the average number of rows and columns per block in the nonuniform test was 256 for each dimension. The number of compute cores in this test set varies from 16 to 240.

Like the tests on the BG/Q system, the data shows an approximately linear increase in computational rate for both uniform and nonuniform tests (see Figures 6 and 7). Unlike the previous tests, the performance difference between the uniform and nonuniform tests is negligibly small; neither set of tests clearly out performs the other. For example, the mean computational times of the fifteen-node, uniform- and nonuniform-matrix tests are 21.1716 and 20.8096 seconds, respectively, with standard deviations of 0.663725 and




Figure 6: Floating operation throughput for multiplication of uniformly- and nonuniformly-blocked square matrices of size 32, 768 on a commodity cluster.




Figure 7: Wall time for multiplication of uniformly- and nonuniformly-blocked square matrices of size 32, 768 on a commodity cluster.

0.544411 seconds, respectively. The mean values of these data points are within one standard deviation of each other, and other data points with the same number of processes are similarly close.

4.3 Efficiency of Task-Based SUMMA

The performance and efficiency of our task-based SUMMA implementation was found to be highly dependent on that of the matrix multiplication kernel, in this case DGEMM (TILEDARRAY supports matrices of standard single- and double-

precision real and complex types as well as user-defined types). The best, single-node performance we achieved with ESSL is 99.1393 GFLOPS, which is 59.7% of Rmax (166.06 GFLOPS per-node) and 48.4% of Rpeak (204.799 GFLOPS per-node). Typically, with a well optimized BLAS library, single-node performance of our task-based SUMMA implementation can achieve 90% to 95% of theoretical peak. For example, the single node performance on our small cluster, using MKL, was 302.787 GFLOPS, which is 90.98% of the theoretical peak performance computed from the base CPU frequency (332.8 GFLOPS per node). The discrepancy in percent of machine peak between these two systems is due to the limited thread scalability of ESSL. Anecdotally, 50% of machine peak is considered “good” performance for ESSL BLAS Level-3 functions with 16 cores.

In Figure 8, we show the percent of efficiency of our algorithm on the IBM BG/Q system, with 100% efficiency set at the maximum measured performance on a single node (99.1393 GFLOPS). The data shows that usually greater




Figure 8: Computational efficiency for multiplication of uniformly- and nonuniformly-blocked square matrices on IBM BG/Q. One hundred percent efficiency refers to the peak, measured performance of a single node, with TILEDARRAY.

than 75% efficiency is maintained across a variety of matrix sizes and number of cores, relative to the single node performance. The efficiency drops down to as low as 50% in the strong scaling limit, when each processor core has 4 or fewer target matrix blocks. The granularity of tasks limits the maximum amount of parallelism achievable for each matrix size, which causes a steep drop in efficiency in the right-most tail of each data set.

4.4 Load Variability with Nonuniform Blocking

The randomly sized matrix blocks in these tests are intended to simulate the irregular blocking structure of data in quantum physics problems. We used a low-quality random number generator to create significant inhomogeneity in

the blocks sizes. The ratio of minimum to maximum memory and work loads for these generated matrices are given in Table 1.

Matrix Size	Memory Min:Max	Work Min:Max
32768	1:2.99	1:4.46
65536	1:3.54	1:5.77
98304	1:3.27	1:5.51
256000	1:3.93	1:7.12

Table 1: Ratio of the minimum to maximum memory and work loads for nonuniformly-blocked matrices.

Despite the substantial inhomogeneity at the block level, because each process hold many blocks due to data overdecomposition, the effective load imbalance in the computation and communication is somewhat smaller than what is reported in Table 1. For example, with a matrix size of 32,768 and 256 processes, we see a ratio of 1:1.35 between processes with the smallest and largest memory usage.

5. CONCLUSIONS

We presented a fine-grained task-based reformulation of Scalable Universal Matrix Multiplication Algorithm (SUMMA). In conjunction with multiple-issue scheduling of SUMMA iterations, the new formulation should be tolerant of data inhomogeneity due to matrix structure. The implementation of our algorithm in TILEDARRAY library performed equally well for multiplication of square matrices with uniform and nonuniform blocking (the latter designed to simulate the domain-specific blocking structure characteristic of the electronic structure domain). The implementation scales well on a small commodity cluster as well as a high-end IBM BG/Q supercomputer, with typical measured efficiencies relative to the single node performance of 75% on as many as 2^{18} cores.

6. ACKNOWLEDGMENTS

This work was supported by the National Science Foundation (grants CHE-1362655 and ACI-1047696), Camille Dreyfus Teacher-Scholar Award, and Alfred P. Sloan Fellowship. We gratefully acknowledge the computer time provided by the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program. This research used resources of the Argonne Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC02-06CH11357. The use of commodity cluster was acquired with the help of the U.S. National Science Foundation award CHE- 0847295. We also acknowledge the started allocation from Advanced Research Computing at Virginia Tech.

We would like to thank Mr. Cannada A. Lewis for the help with data analysis and useful contributions to TILEDARRAY.

7. REFERENCES

- [1] AGARWAL, R. C., BALLE, S. M., GUSTAVSON, F. G., JOSHI, M., AND PALKAR, P. A three-dimensional approach to parallel matrix multiplication. *IBM J. Res. & Dev.* 39, 5 (Sept. 1995), 575–582.
- [2] AGARWAL, R. C., GUSTAVSON, F. G., AND ZUBAIR, M. A high-performance matrix-multiplication

- algorithm on a distributed-memory parallel computer, using overlapped communication. *IBM J. Res. & Dev.* 38, 6 (Nov. 1995), 673–681.
- [3] BALLARD, G., BULUÇ, A., DEMMEL, J., GRIGORI, L., LIPSHITZ, B., SCHWARTZ, O., AND TOLEDO, S. Communication optimal parallel multiplication of sparse random matrices. In *SPAA '13: Proceedings of the twenty-fifth annual ACM symposium on Parallelism in algorithms and architectures* (July 2013), ACM Request Permissions.
- [4] BALLARD, G., DEMMEL, J., HOLTZ, O., LIPSHITZ, B., AND SCHWARTZ, O. Brief announcement: strong scaling of matrix multiplication algorithms and memory-independent communication lower bounds. In *SPAA '12* (New York, New York, USA, 2012), ACM Press, pp. 77–79.
- [5] BALLARD, G., DEMMEL, J., HOLTZ, O., LIPSHITZ, B., AND SCHWARTZ, O. Communication-optimal parallel algorithm for strassen's matrix multiplication. In *SPAA '12* (New York, New York, USA, 2012), ACM Press, p. 193.
- [6] BARUCH, E. *Designing Communication-efficient Matrix Algorithms in Distributed-memory Cilk*. Tel-Aviv University, 2001.
- [7] BAUMGARTNER, G., BERNHOLDT, D. E., COCIORVA, D., HARRISON, R. J., HIRATA, S., LAM, C.-C., NOOLJEN, M., PITZER, R. M., RAMANUJAM, J., AND SADAYAPPAN, P. A high-level approach to synthesis of high-performance codes for quantum chemistry. pp. 1–10.
- [8] BERNTSEN, J. Communication efficient matrix multiplication on hypercubes. *Parallel Computing* 12, 3 (Dec. 1989), 335–342.
- [9] BOSILCA, G., BOUTELLER, A., DANALIS, A., FAVERGE, M., HAIDAR, A., HERAULT, T., KURZAK, J., LANGOU, J., LEMARINIER, P., LTAIEF, H., LUSZCZEK, P., YARKHAN, A., AND DONGARRA, J. Flexible development of dense linear algebra algorithms on massively parallel architectures with DPLASMA. In *Proceedings of PDSEC 2011* (Anchorage, United States, may 2011).
- [10] BULUÇ, A., AND GILBERT, J. R. Challenges and advances in parallel sparse matrix-matrix multiplication. In *Proceedings of the 37th Int'l Conf. on Par. Proc. (ICPP '08)* (2008), pp. 503–510.
- [11] BULUÇ, A., AND GILBERT, J. R. Parallel Sparse Matrix-Matrix Multiplication and Indexing: Implementation and Experiments. *SIAM J. Sci. Comput.* 34, 4 (Jan. 2012), C170–C191.
- [12] CALVIN, J. A., AND VALEEV, E. F. Tiledarray: A massively-parallel, block-sparse tensor library written in c++. <https://github.com/valeevgroup/tiledarray/>, 2015.
- [13] CANNON, L. E. *A cellular computer to implement the kalman filter algorithm*. PhD thesis, Ann Arbor, July 1969.
- [14] CHOI, J. A fast scalable universal matrix multiplication algorithm on distributed-memory concurrent computers. In *IPPS-97* (Apr. 1997), IEEE Comput. Soc. Press, pp. 310–314.
- [15] CHOI, J. A new parallel matrix multiplication algorithm on distributed-memory concurrent computers. *Concurr. Comput.* 10, 8 (July 1998), 655–670.
- [16] CHOI, J., DONGARRA, J. J., AND WALKER, D. W. PB-BLAS: a set of parallel block basic linear algebra subprograms. *Concurrency: Pract. Exper.* 8, 7 (Sept. 1996), 517–535.
- [17] CHOI, J., WALKER, D. W., AND DONGARRA, J. J. Puma: Parallel universal matrix multiplication algorithms on distributed memory concurrent computers. *Concurr. Comput.* 6, 7 (Oct. 1994), 543–570.
- [18] COHN, H., AND UMANS, C. A group-theoretic approach to fast matrix multiplication. In *SFCS-03* (2003), IEEE Computer. Soc, pp. 438–449.
- [19] DEMMEL, J., ELIAHU, D., FOX, A., KAMIL, S., LIPSHITZ, B., SCHWARTZ, O., AND SPILLINGER, O. Communication-optimal parallel recursive rectangular matrix multiplication. In *2013 IEEE International Symposium on Parallel & Distributed Processing (IPDPS)* (2013), IEEE, pp. 261–272.
- [20] EPIFANOVSKY, E., WORMIT, M., KUŠ, T., LANDAU, A., ZUEV, D., KHISTYAEV, K., MANOHAR, P., KALIMAN, I., DREUW, A., AND KRYLOV, A. I. New implementation of high-level correlated methods using a general block tensor library for high-performance electronic structure calculations. *J Comput Chem* 34, 26 (July 2013), 2293–2309.
- [21] FOX, G. C. *Solving problems on concurrent processors*. Prentice Hall, Englewood Cliffs, N.J, 1988.
- [22] FOX, G. C., OTTO, S. W., AND HEY, A. J. G. Matrix algorithms on a hypercube I: Matrix multiplication. *Parallel Comput.* 4, 1 (1987), 17–31.
- [23] GUNNELS, J., LIN, C., MORROW, G., AND VAN DE GEIJN, R. A. A flexible class of parallel matrix multiplication algorithms. In *Parallel Processing Symposium, 1998. IPPS/SPDP 1998. Proceedings of the First Merged International ... and Symposium on Parallel and Distributed Processing 1998* (1998), pp. 110–116.
- [24] HARRISON, R., AND MADNESS COMMUNITY. MADNESS project hosted on Github. <https://github.com/m-a-d-n-e-s-s/madness>, 2015.
- [25] HIRATA, S. Tensor Contraction Engine: Abstraction and automated parallel implementation of configuration-interaction, coupled-cluster, and many-body perturbation theories. *J. Phys. Chem. A* 107, 46 (Nov. 2003), 9887–9897.
- [26] HOLLMAN, D. S., SCHAEFER, H. F., AND VALEEV, E. F. A tight distance-dependent estimator for screening three-center Coulomb integrals over Gaussian basis functions. *J Chem Phys* 142, 15 (2015), 154106, arXiv:chem-ph/1410.4882.
- [27] HUSBANDS, P., AND YELICK, K. Multi-threading and one-sided communication in parallel LU factorization. In *SC '07* (New York, New York, USA, 2007), ACM Press, p. 1.
- [28] HUSS-LEDERMAN, S., JACOBSON, E. M., AND TSAO, A. Comparison of scalable parallel matrix multiplication libraries. In *SPLC-93* (1993), IEEE Comput. Soc. Press, pp. 142–149.
- [29] HUTCHINSON, B., DENG, L., AND YU, D. Tensor Deep Stacking Networks. *IEEE Trans. Pattern Anal.*

- Mach. Intell.* 35, 8 (2012), 1944–1957.
- [30] IRONY, D., TOLEDO, S., AND TISKIN, A. Communication lower bounds for distributed-memory matrix multiplication. *Journal of Parallel and Distributed Computing* 64, 9 (Sept. 2004), 1017–1026.
 - [31] LAI, P.-W., STOCK, K., RAJBHANDARI, S., KRISHNAMOORTHY, S., AND SADAYAPPAN, P. A framework for load balancing of tensor contraction expressions via dynamic task partitioning. In *SC '13: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis* (Nov. 2013), ACM Request Permissions.
 - [32] LIN, C., AND SNYDER, L. A matrix product algorithm and its comparative performance on hypercubes. In *Scalable High Performance Computing Conference SHPCC-92*. (1992), IEEE Comput. Soc. Press, pp. 190–194.
 - [33] MØRUP, M., HANSEN, L. K., ARNFRED, S. M., LIM, L.-H., AND MADSEN, K. H. Shift-invariant multilinear decomposition of neuroimaging data. *NeuroImage* 42, 4 (Oct. 2008), 1439–1450.
 - [34] NEESE, F., HANSEN, A., AND LIAKOS, D. G. Efficient and accurate approximations to the local coupled cluster singles doubles method using a truncated pair natural orbital basis. *J. Chem. Phys.* 131, 6 (Aug. 2009), –.
 - [35] NEESE, F., WENNMOHS, F., AND HANSEN, A. Efficient and accurate local approximations to coupled-electron pair approaches: An attempt to revive the pair natural orbital method. *J. Chem. Phys.* 130, 11 (2009), 1–18.
 - [36] NGUYEN, T., CICOTTI, P., BYLASKA, E., QUINLAN, D., AND BADEN, S. B. Bamboo: Translating MPI applications to a latency-tolerant, data-driven form. In *SC '12: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis* (Nov. 2012), IEEE Computer Society Press.
 - [37] OZOG, D., SHENDE, S., MALONY, A., HAMMOND, J. R., DINAN, J., AND BALAJI, P. Inspector/executor load balancing algorithms for block-sparse tensor contractions. In *ICS '13: Proceedings of the 27th international ACM conference on International conference on supercomputing* (June 2013), ACM.
 - [38] PAVOŠEVIĆ, F., NEESE, F., AND VALEEV, E. F. Geminal-spanning orbitals make explicitly correlated reduced-scaling coupled-cluster methods robust, yet simple. *J. Chem. Phys.* 141, 5 (Aug. 2014), 1–10.
 - [39] POULSON, J., MARKER, B., HAMMOND, J. R., ROMERO, N. A., AND VAN DE GEIJN, R. A. Elemental: A new framework for distributed memory dense matrix computations. *ACM Trans. Math. Softw.* (June 2010), 1–26.
 - [40] RAJBHANDARI, S., NIKAM, A., LAI, P.-W., STOCK, K., KRISHNAMOORTHY, S., AND SADAYAPPAN, P. A communication-optimal framework for contracting distributed tensors. In *SC '14: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (Nov. 2014), IEEE Press.
 - [41] RAJBHANDARI, S., NIKAM, A., LAI, P.-W., STOCK, K., KRISHNAMOORTHY, S., AND SADAYAPPAN, P. CAST: Contraction algorithm for symmetric tensors. In *Parallel Processing (ICPP), 2014 43rd International Conference on* (2014), IEEE, pp. 261–272.
 - [42] RUBENSSON, E. H., AND RUDBERG, E. Locality-aware parallel block-sparse matrix-matrix multiplication using the Chunks and Tasks programming model. 1–17, 1501.07800.
 - [43] SCHATZ, M. D., POULSON, J., AND VAN DE GEIJN, R. A. Scalable universal matrix multiplication algorithms: 2D and 3D variations on a theme. *submitted to ACM Transactions on Mathematical Software* (2012), 1–30.
 - [44] SMILDE, A., BRO, R., AND GELADI, P. *Multi-way Analysis: Applications in the Chemical Sciences*. Wiley, Chichester, Oct. 2004.
 - [45] SOLOMONIK, E., BHATELE, A., AND DEMMEL, J. Improving communication performance in dense linear algebra via topology aware collectives. In *SC '11* (New York, New York, USA, 2011), ACM Press, p. 1.
 - [46] SOLOMONIK, E., AND DEMMEL, J. Communication-optimal parallel 2.5D matrix multiplication and LU factorization algorithms. In *Euro-Par 2011 Parallel Processing*. Springer Berlin Heidelberg, 2011, pp. 90–109.
 - [47] SOLOMONIK, E., MATTHEWS, D., HAMMOND, J. R., STANTON, J. F., AND DEMMEL, J. A massively parallel tensor contraction framework for coupled-cluster computations. *J. Parallel Distrib. Comput.* 74, 12 (Dec. 2014), 3176–3190.
 - [48] STRASSEN, V. Gaussian elimination is not optimal. *Numer. Math.* 13, 4 (Aug. 1969), 354–356.
 - [49] TAN, L., CHEN, L., CHEN, Z., ZONG, Z., LI, D., AND GE, R. Improving performance and energy efficiency of matrix multiplication via pipeline broadcast. In *Cluster Computing (CLUSTER), 2013 IEEE International Conference on* (Sept. 2013), IEEE, pp. 1–5.
 - [50] VAN DE GEIJN, R. A., AND WATTS, J. SUMMA: Scalable universal matrix multiplication algorithm. *Concurr. Comput.* 9, 4 (1997), 255–274.